CS 3630!

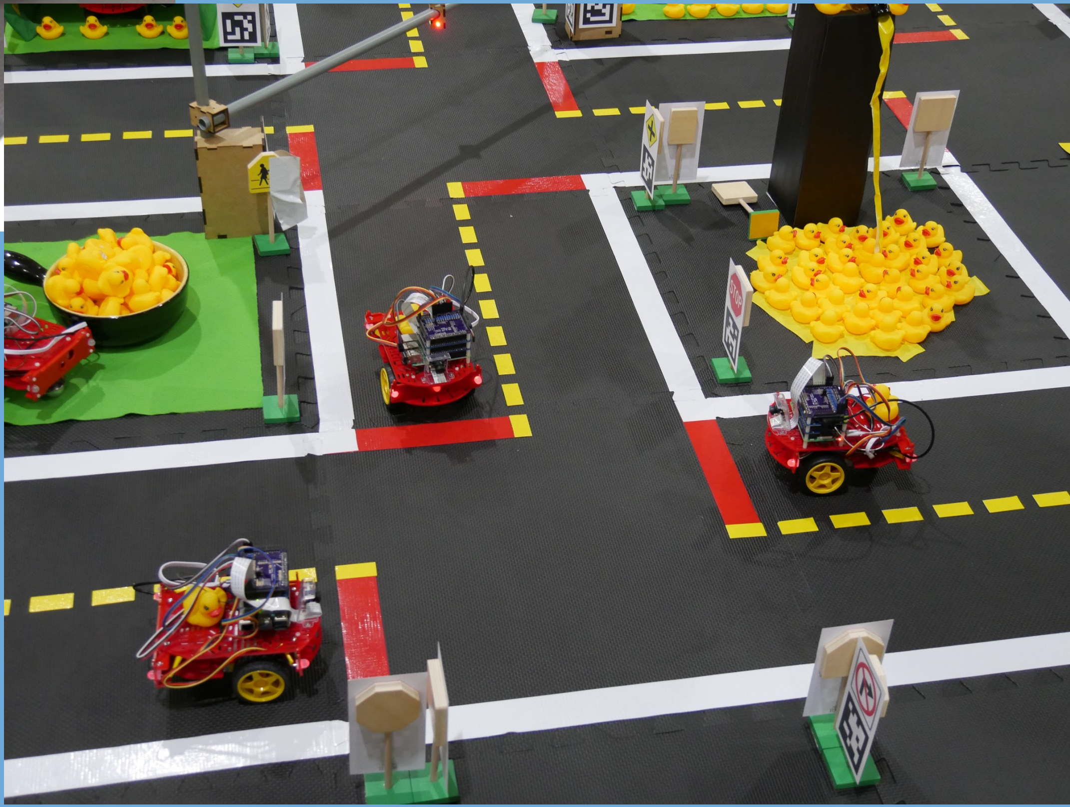

Lecture 19:

Deep Learning in Robotics

Images and text sampled from a selection of 2019 research papers.

1. Applications in Perception

- From pixels to concepts:
    - Image processing
    - Object classification
    - Object detection
    - Pixelwise segmentation

Colorization

- Given a grayscale image, colorize the image realistically
- Zhang et al. pose colorization as classification task and use class-rebalancing to improve results
- Demonstrate higher rates of fooling humans using "colorization Turing test"

Colorful Image Colorization. Richard Zhang, Phillip Isola, Alexei A. Efros. ECCV 2016.

DeOldify

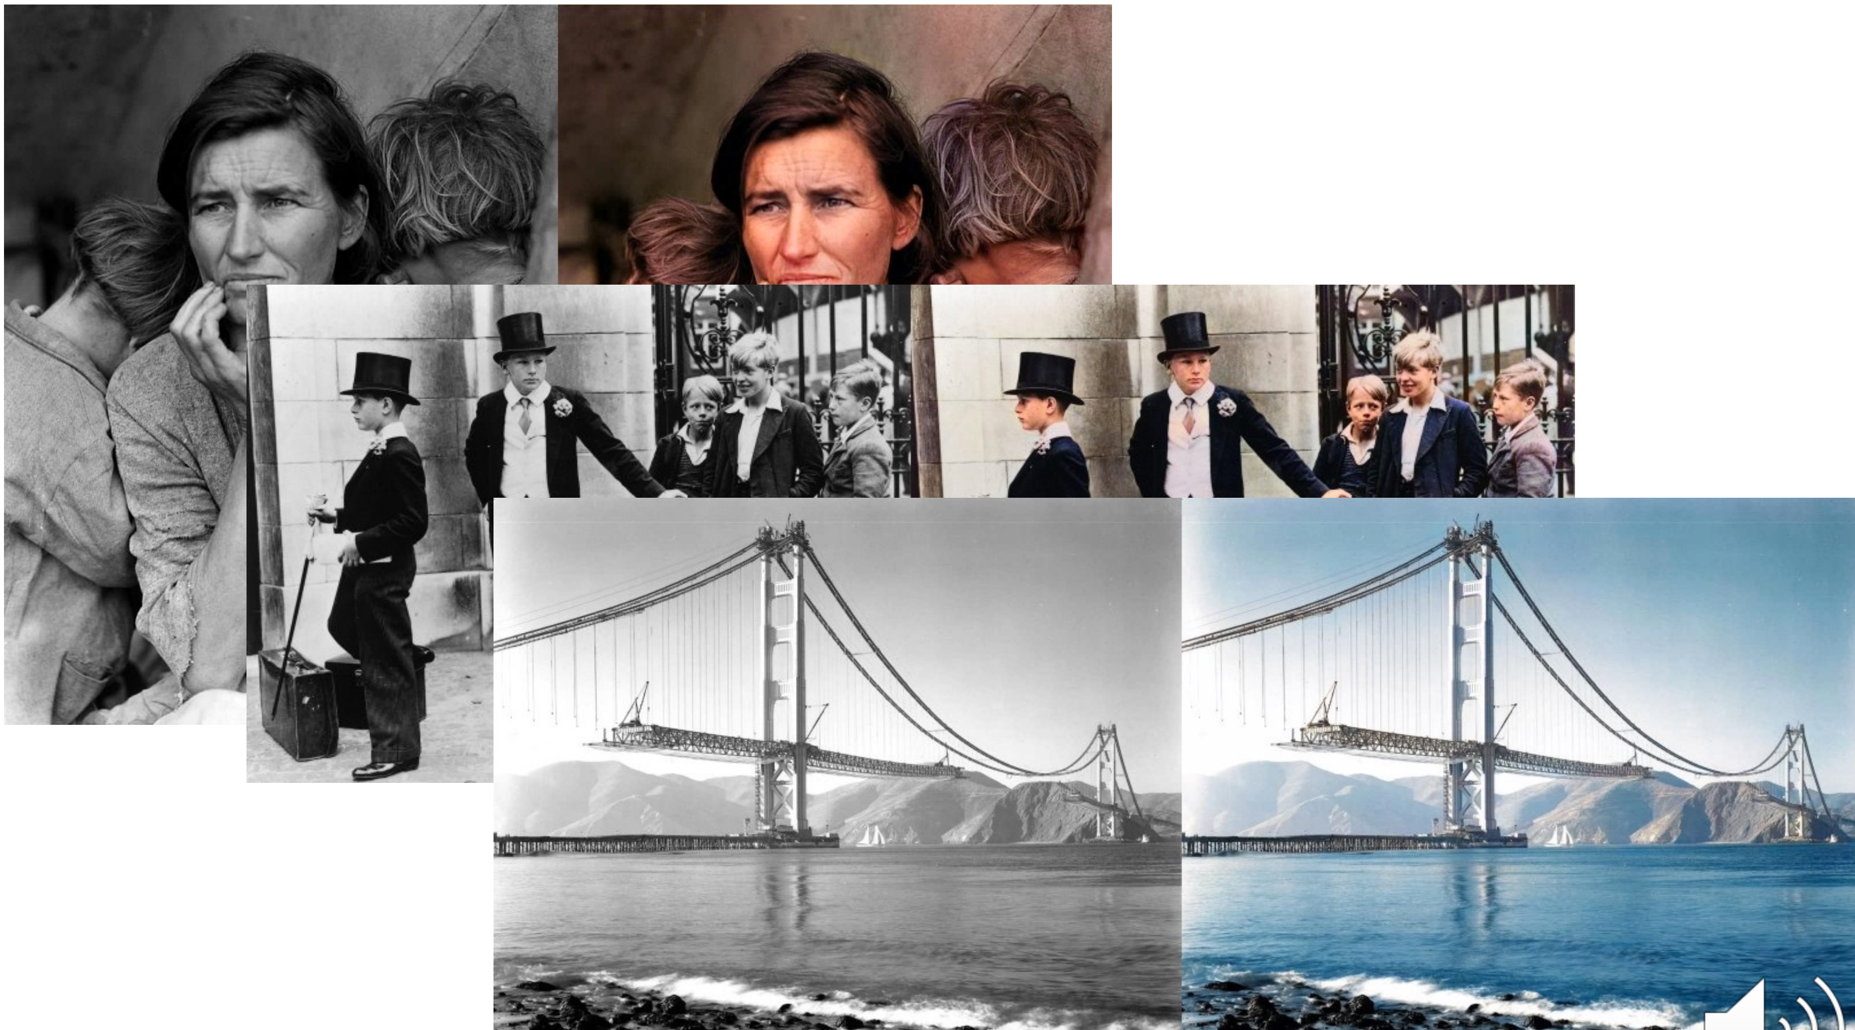

https://github.com/jantic/DeOldify

Super-Resolution

Low resolution

High resolution

Object Classification Revolution

Results: ILSVRC 2012

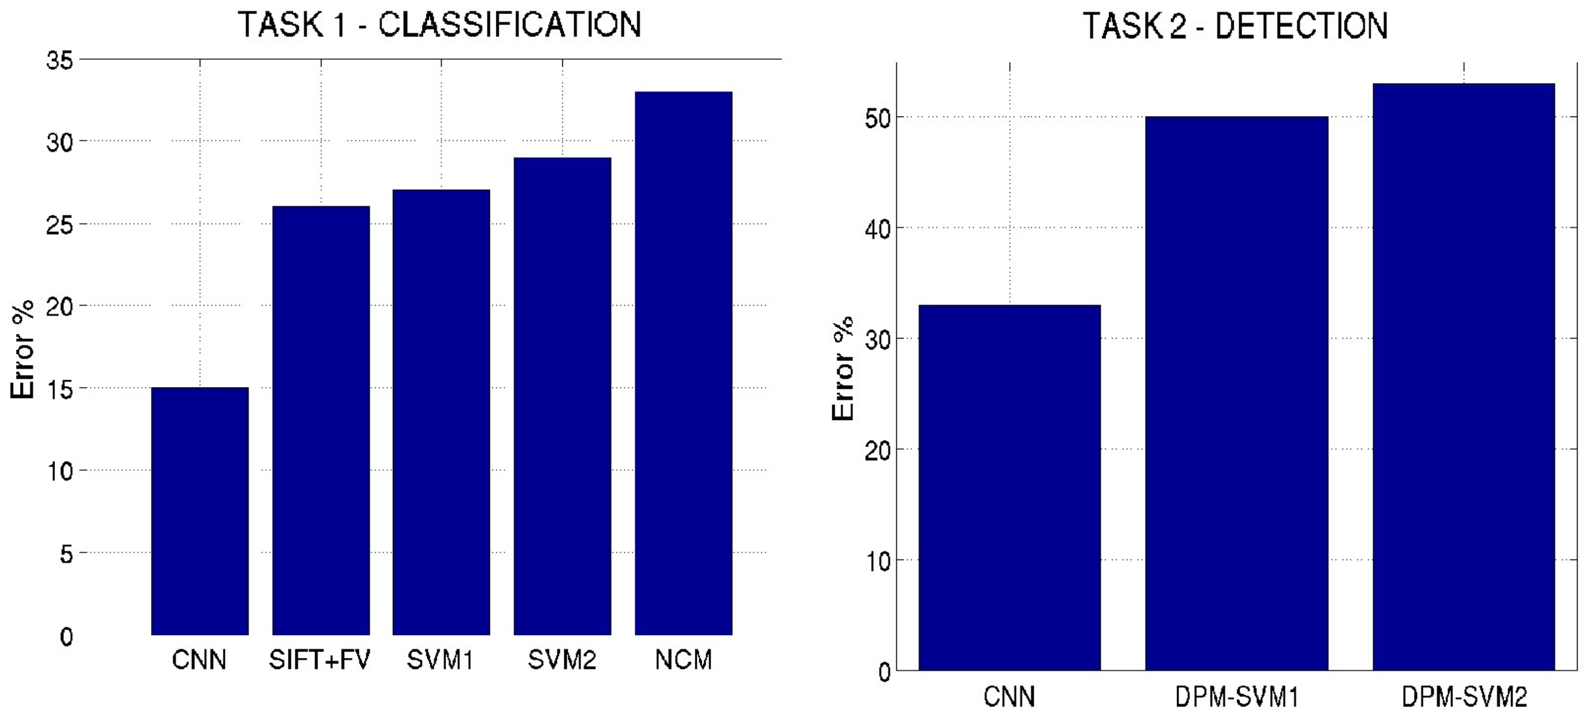

Krizhevsky et al. "ImageNet Classification with deep CNNs" NIPS 2012

Ranzato

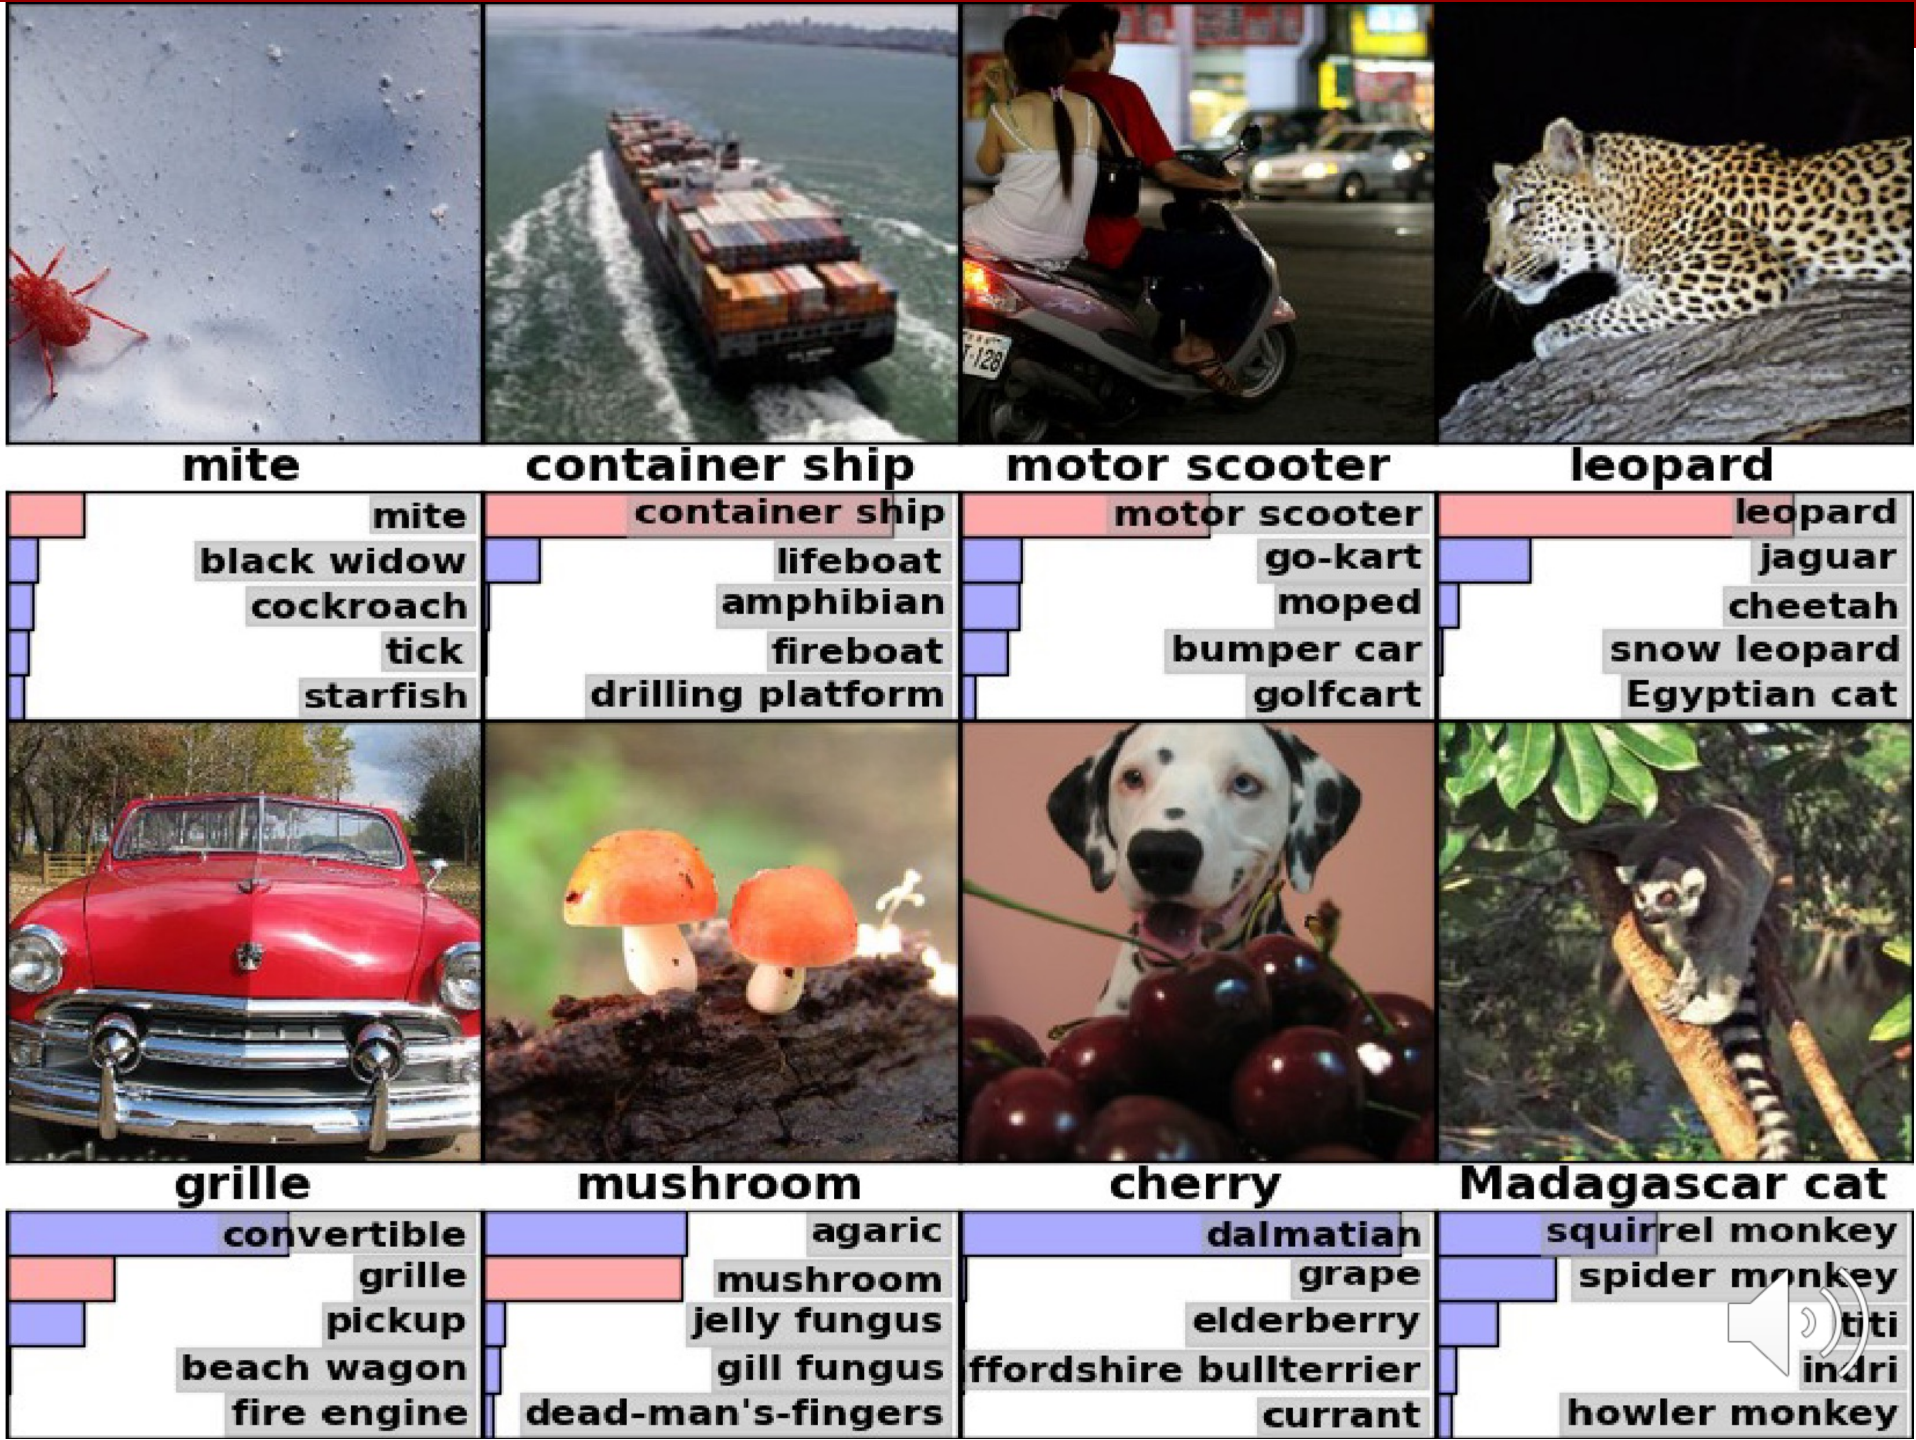

mite
mite
black widow
cockroach
tick
starfish
container ship
container ship
lifeboat
amphibian
fireboat
drilling platform
motor scooter
motor scooter
go-kart
moped
bumper car
golfcart
leopard
leopard
jaguar
cheetah
snow leopard
Egyptian cat
grille
convertible
grille
pickup
beach wagon
fire engine
mushroom
agaric
mushroom
jelly fungus
gill fungus
dead-man's-fingers
cherry
dalmatian
grape
elderberry
ffordshire bullterrier
currant
Madagascar cat
squirrel monkey
spider monkey
titi
indri
howler monkey

Revolution of Depth

AlexNet, 8 layers

(ILSVRC 2012)

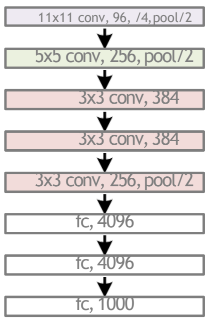

VGG, 19 layers

(ILSVRC 2014)

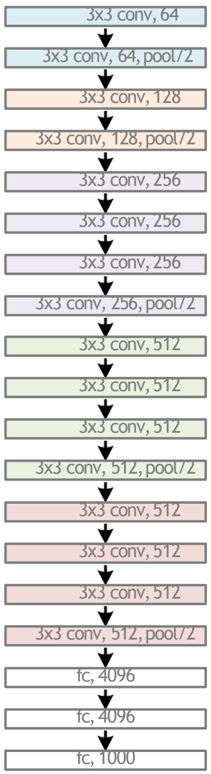

GoogleNet, 22 layers

(ILSVRC 2014)

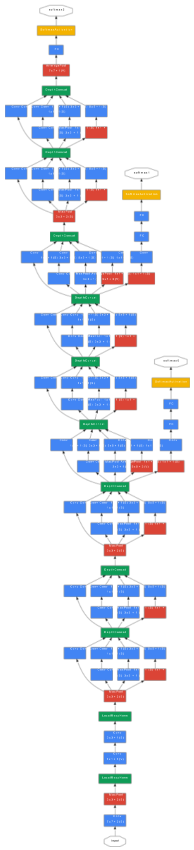

Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. “Deep Residual Learning for Image Recognition”. CVPR 2016.

Object Detection

Image Classification (what?)

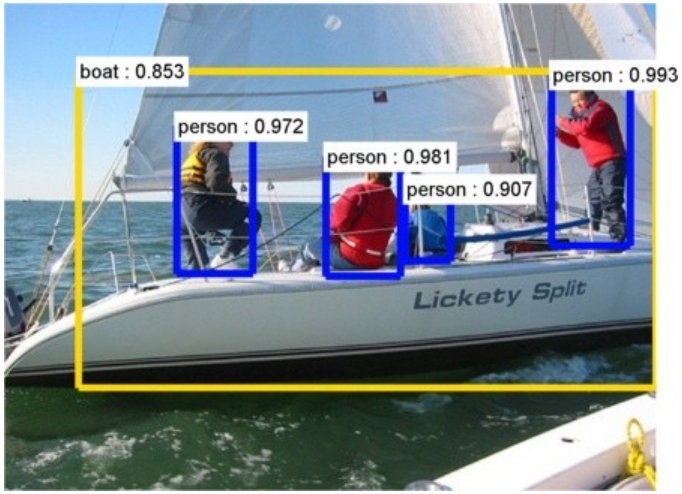

Object Detection (what + where?)

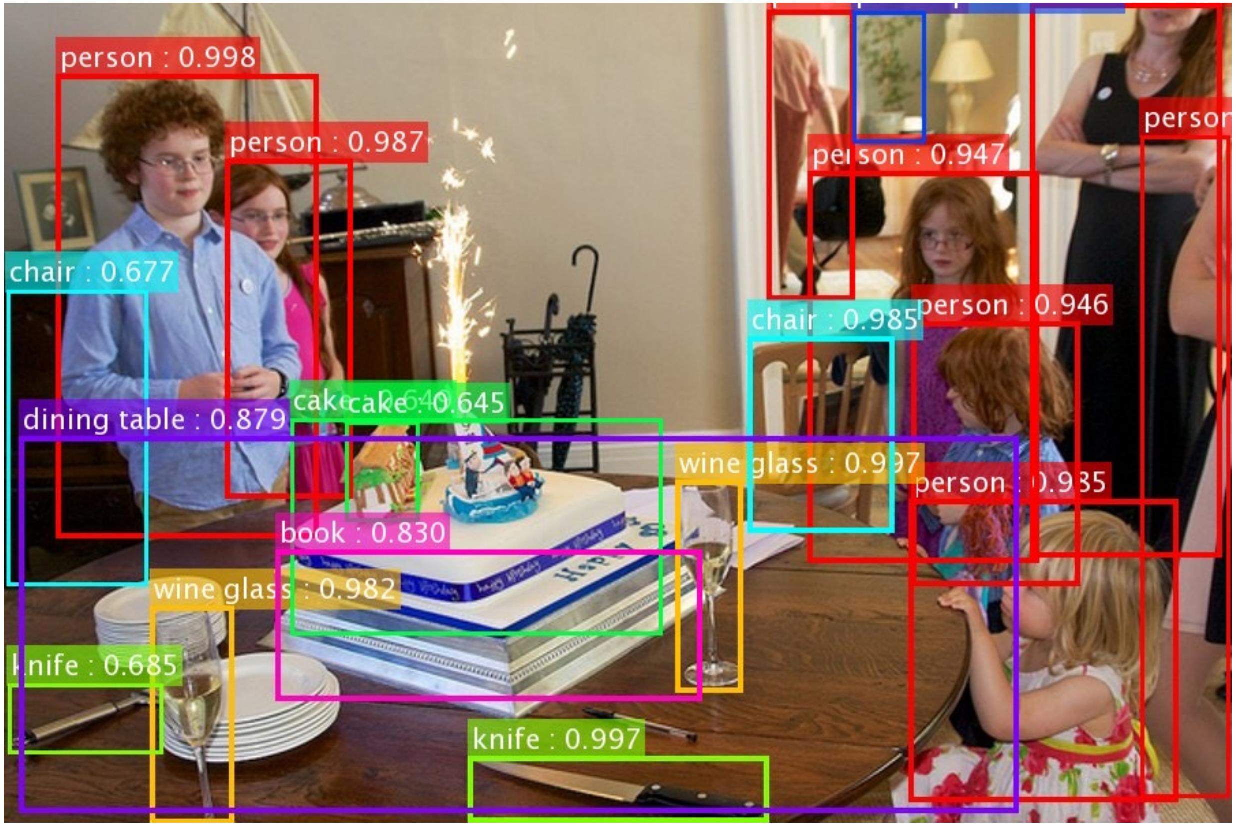

ResNet’s object detection result on COCO

Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. “Deep Residual Learning for Image Recognition”. CVPR 2016.
Shaoqing Ren, Kaiming He, Ross Girshick, & Jian Sun. “Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks”. NIPS 2015.

this video is available online: https://youtu.be/WZmSMkK9VuA

Results on real video. Models trained on MS COCO (80 categories). (frame-by-frame; no temporal processing)

Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. "Deep Residual Learning for Image Recognition". CVPR 2016.
Shaoqing Ren, Kaiming He, Ross Girshick, & Jian Sun. "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks". NIPS 2015.

Detectron

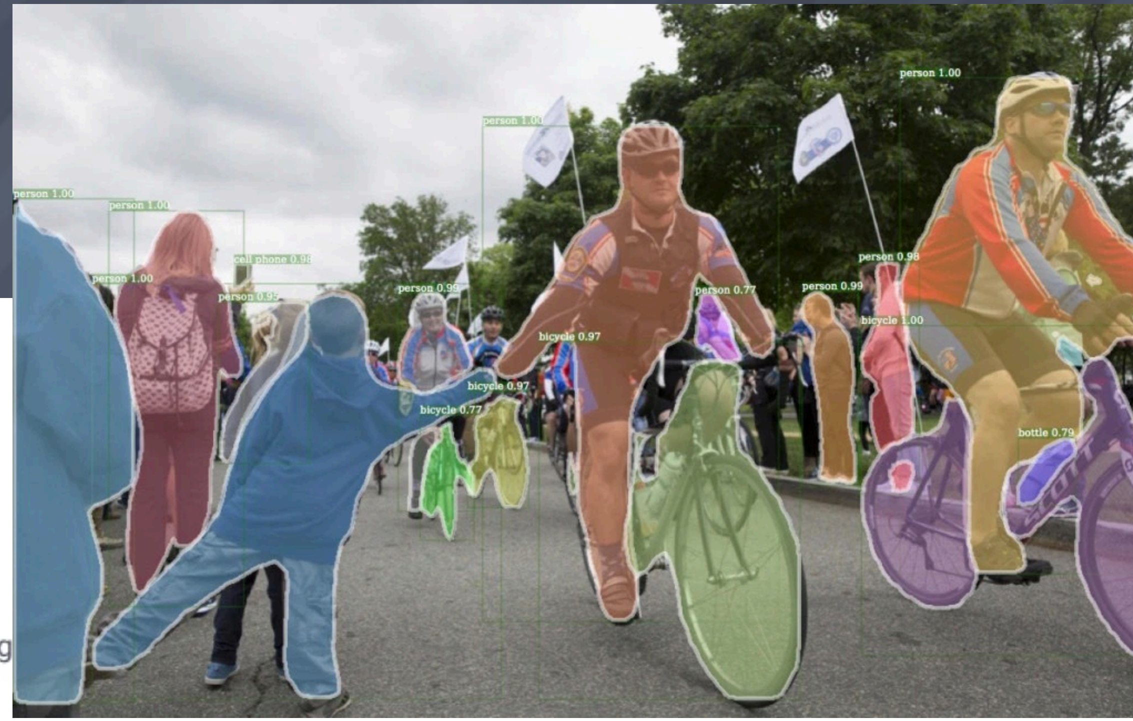

Detectron includes implementations of the following

- Mask R-CNN — *Marr Prize at ICCV 2017*
- RetinaNet — *Best Student Paper Award at ICCV 2017*
- Faster R-CNN
- RPN
- Fast R-CNN
- R-FCN

using the following backbone network architectures:

- ResNeXt{50,101,152}
- ResNet{50,101,152}
- Feature Pyramid Networks (with ResNet/ResNeXt)
- VGG16

2. Applications in Robotics

- From pixels to action:
  - Deep Stereo
  - Depth from a Single Image
  - Best papers from ICRA 2019
  - Best papers from CorL 2019

Deep Stereo

Computing the Stereo Matching Cost with a Convolutional Neural Network

Jure Žbontar
University of Ljubljana
jure.zbontar@fri.uni-lj.si

Yann LeCun
New York University
yann@cs.nyu.edu

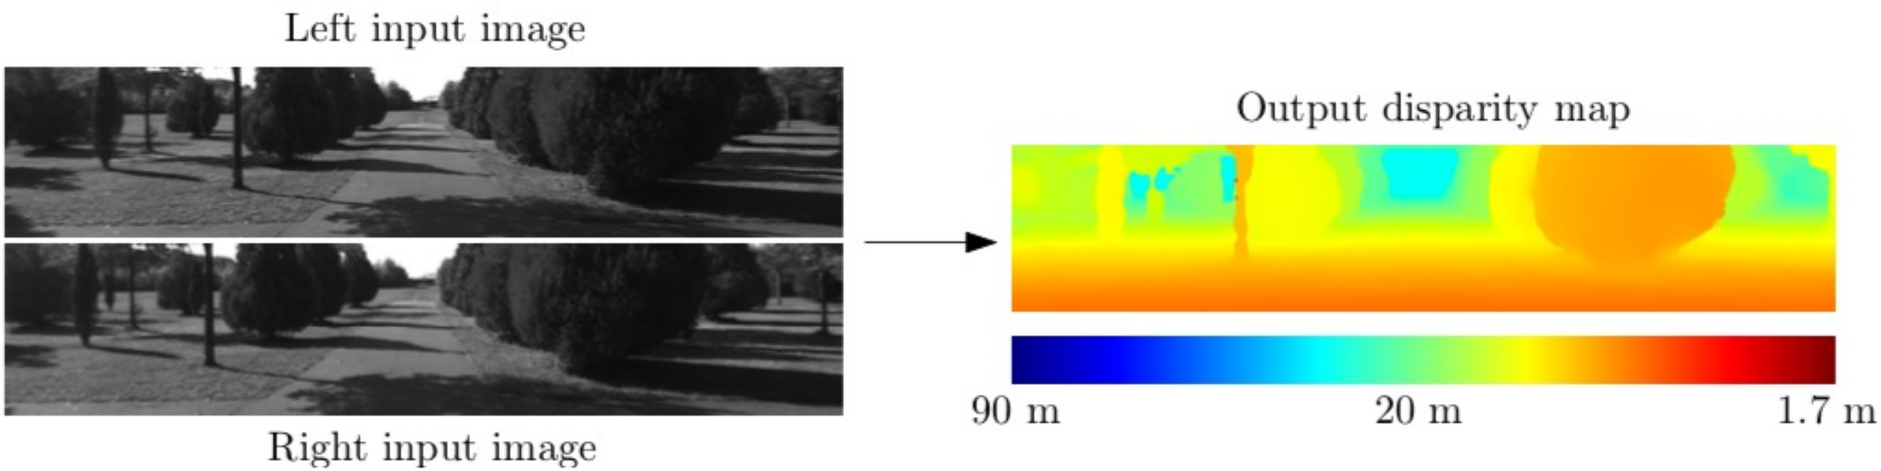

- Learns cost function

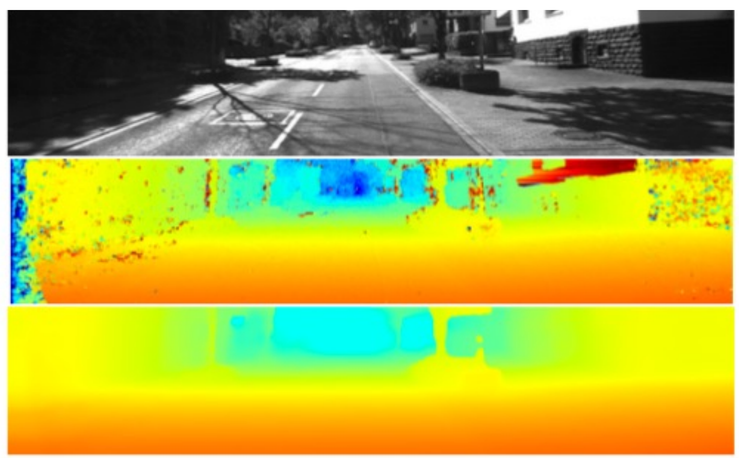

Winner-take all

Smoothed

Stereo Datasets

- *FlyingThings3D -> DispNet*

A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation

Nikolaus Mayer*[1], Eddy Ilg*[1], Philip Häusser*[2], Philipp Fischer*[1†]
[1]University of Freiburg [2]Technical University of Munich
[1]{mayern,ilg,fischer}@cs.uni-freiburg.de [2]haeusser@cs.tum.edu
Daniel Cremers
Technical University of Munich
cremers@tum.de
Alexey Dosovitskiy, Thomas Brox
University of Freiburg
{dosovits,brox}@cs.uni-freiburg.de

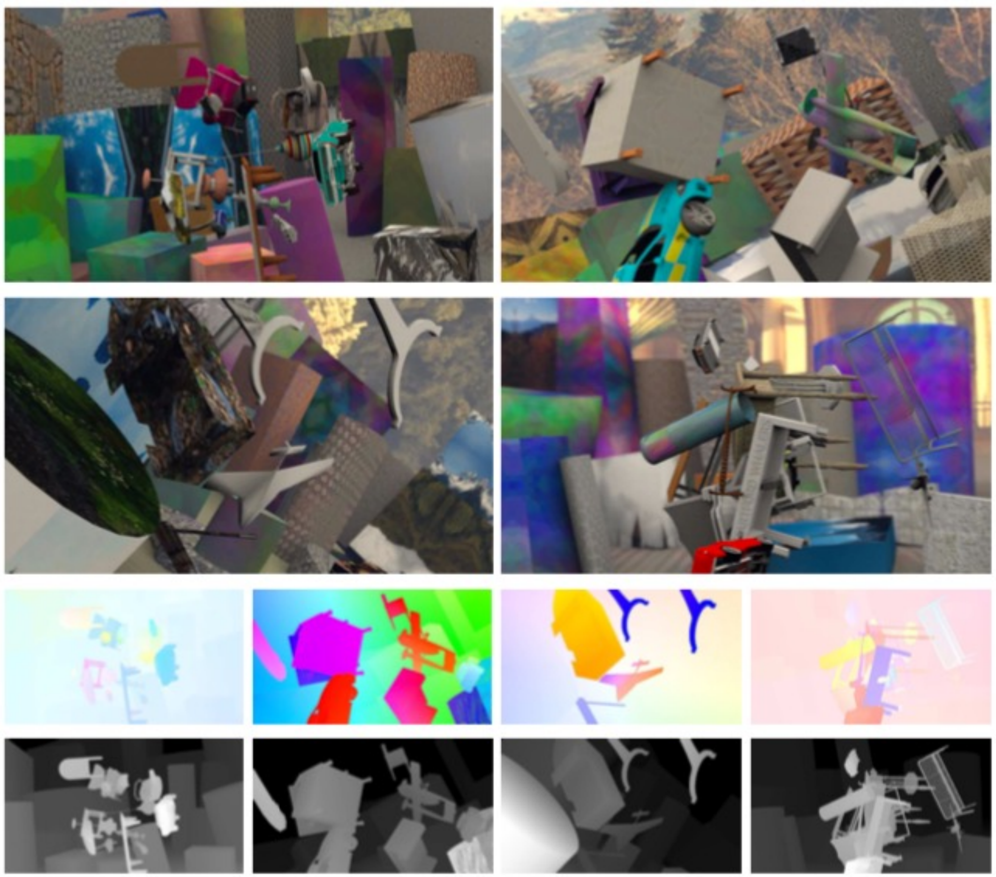

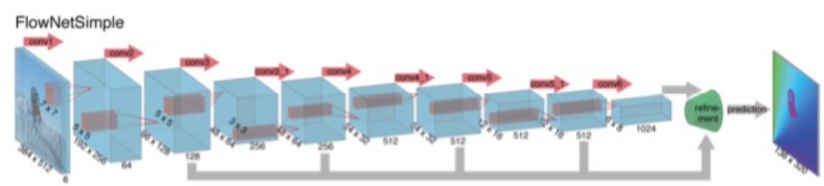

disp GT / prediction

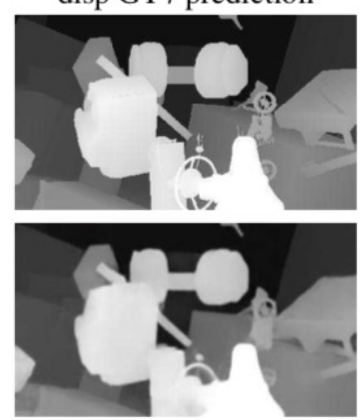

Unsupervised Monocular Depth Estimation with Left-Right Consistency

Clément Godard Oisin Mac Aodha Gabriel J. Brostow
University College London
http://visual.cs.ucl.ac.uk/pubs/monoDepth/

Monocular Depth

- Can we learn depth from a *single* image?
- Train on stereo, but test on mono!
- Learn to war left to right and vice versa

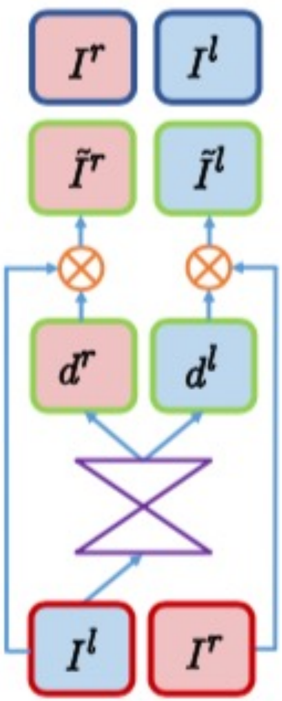

Ours

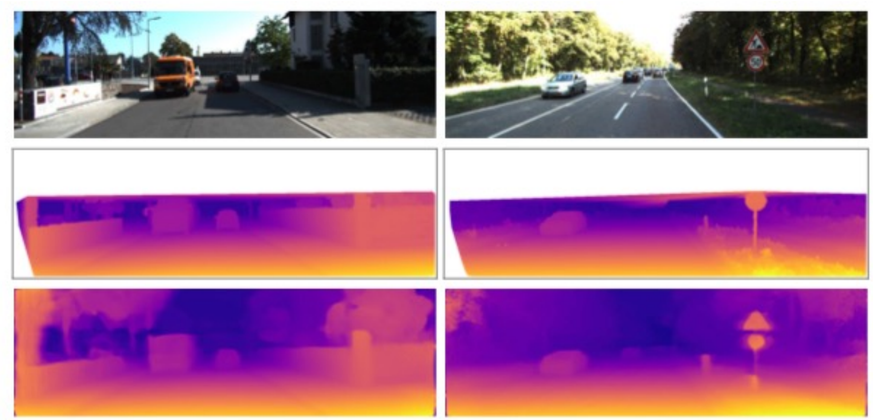

Figure 1. Our depth prediction results on KITTI 2015. Top to bottom: input image, ground truth disparities, and our result. Our method is able to estimate depth for thin structures such as street signs and poles.

ICRA 2019

The International Conference on Robotics and Automation (May 20-24) is the flagship conference of the IEEE Robotics and Automation Society, bringing together the world's top researchers and companies to share ideas and advances in the field.

SUBJECTS

All

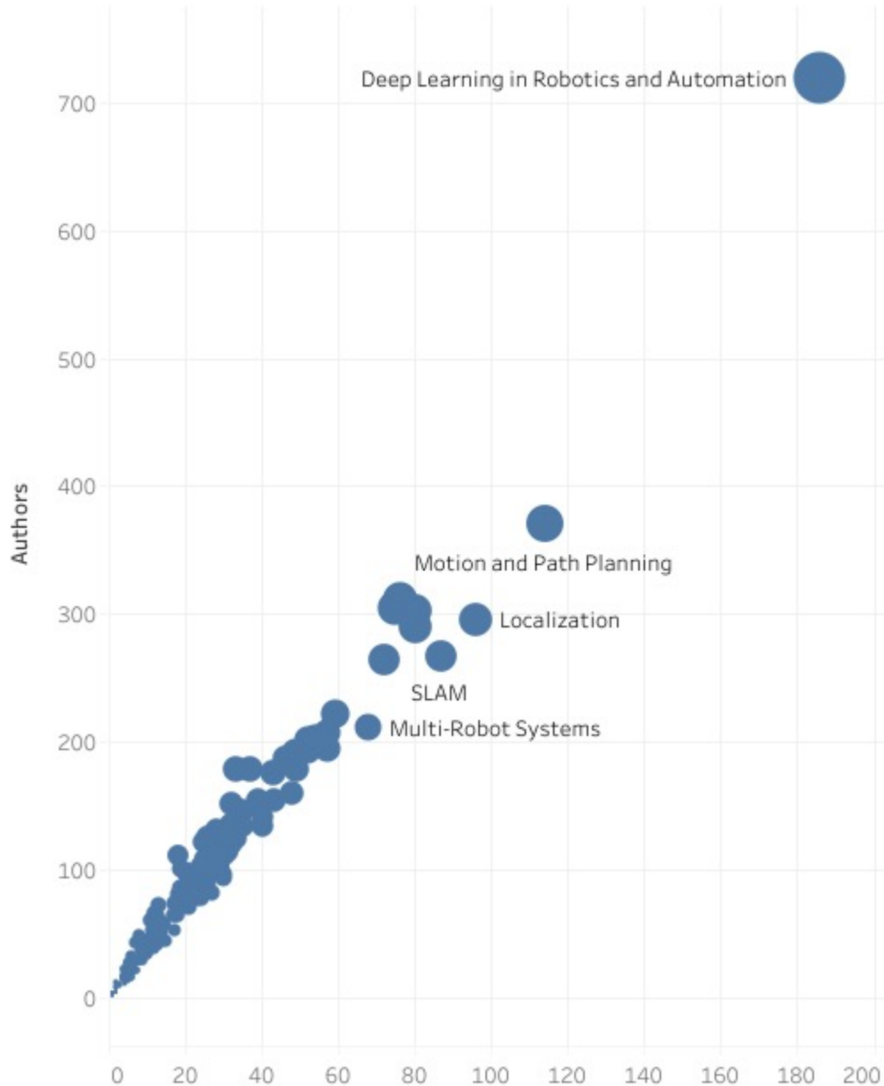

PICK INSTITUTION

(All)

◀ Clear filter;
Uncheck "All";
Type in Selections

INSTITUTION OVERVIEW

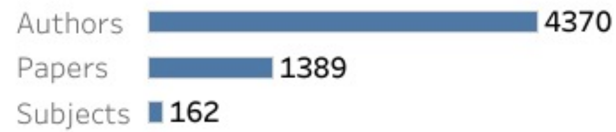

TOP 10 SUBJECTS Papers | Authors

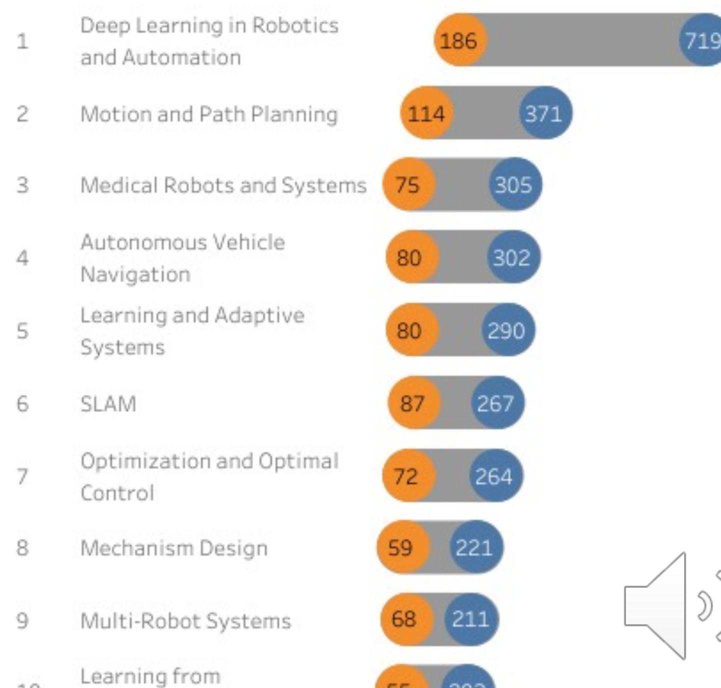

Localization in driving (best paper runner up)

Variational End-to-End Navigation and Localization

Alexander Amini[1], Guy Rosman[2], Sertac Karaman[3] and Daniela Rus[1]

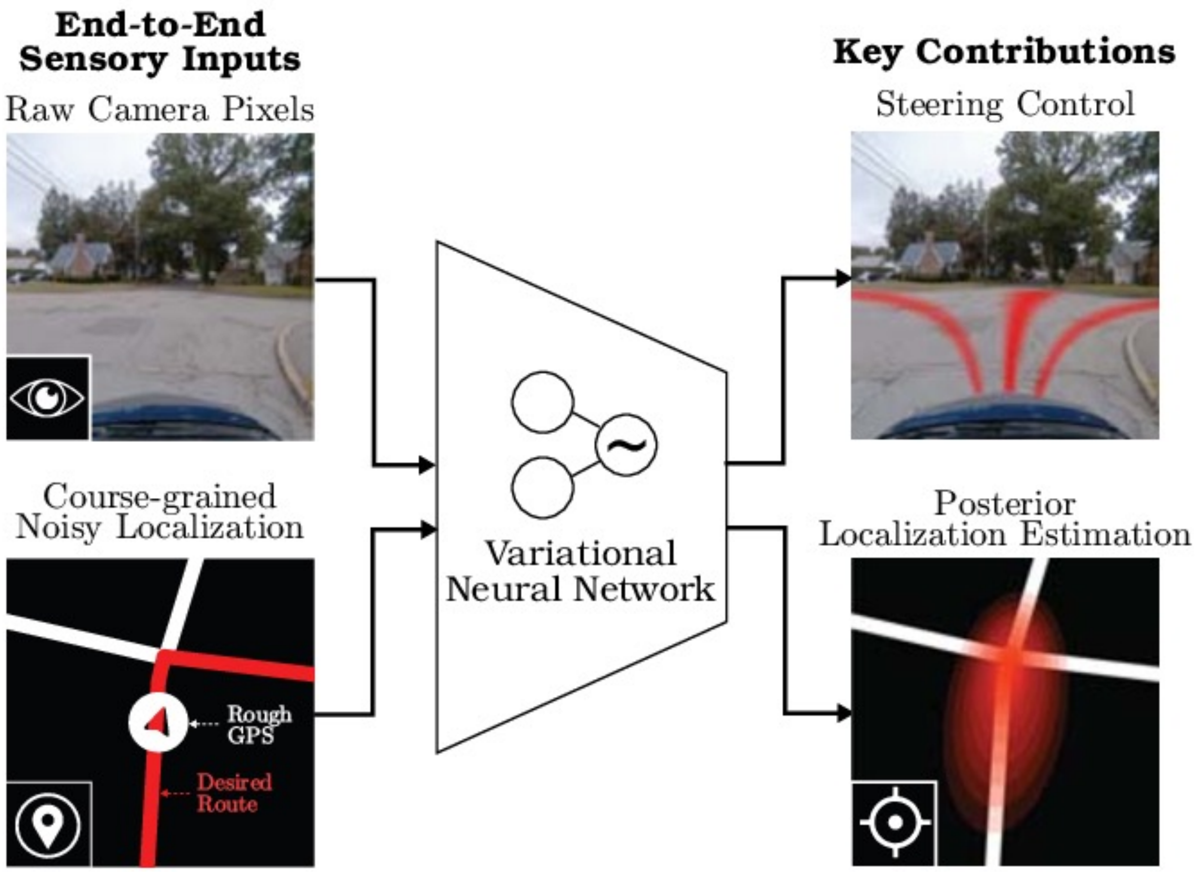

Estimating tactile properties from images (2nd runner up)

Deep Visuo-Tactile Learning:
Estimation of Tactile Properties from Images

Kuniyuki Takahashi, Jethro Tan

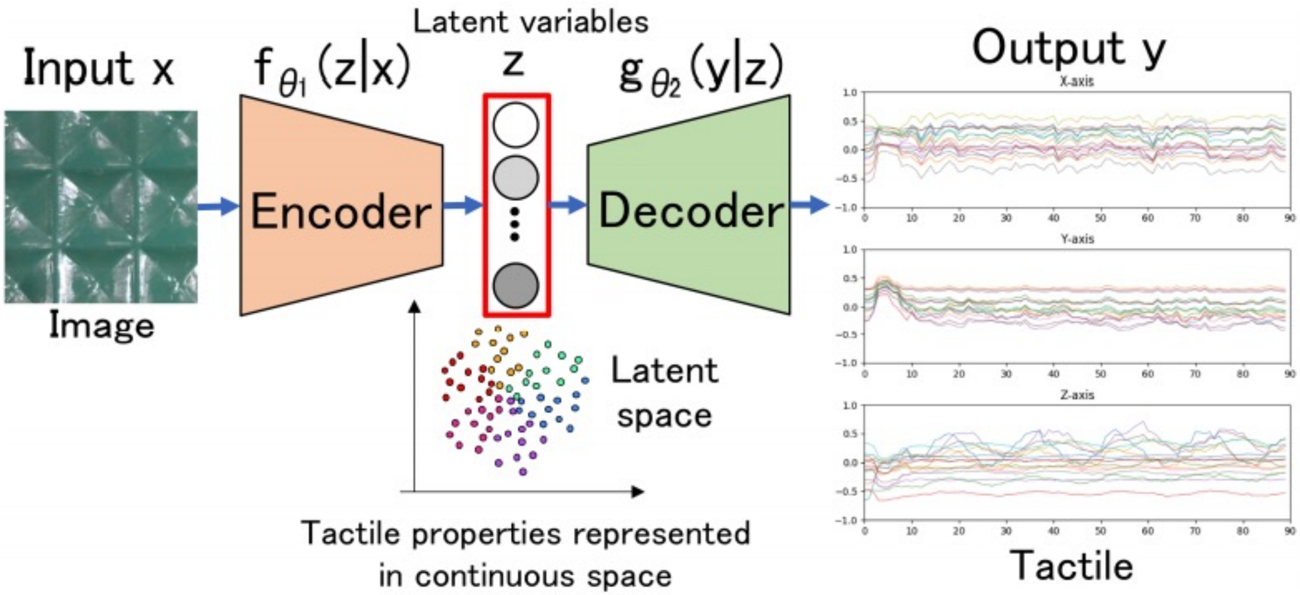

Fig. 2: Proposed network architecture for deep visuo-tactile learning composed of encoder-decoder layers and latent variables. Input is texture image of material and, output is the tactile data contains measured forces by a tactile sensor in the x, y, and z axes. After training, latent variables would contain tactile properties of materials correlating images with tactile sense.

Multimodal perception (best paper)

Making Sense of Vision and Touch: Self-Supervised Learning of Multimodal Representations for Contact-Rich Tasks

Michelle A. Lee*, Yuke Zhu*, Krishnan Srinivasan, Parth Shah,
Silvio Savarese, Li Fei-Fei, Animesh Garg, Jeannette Bohg

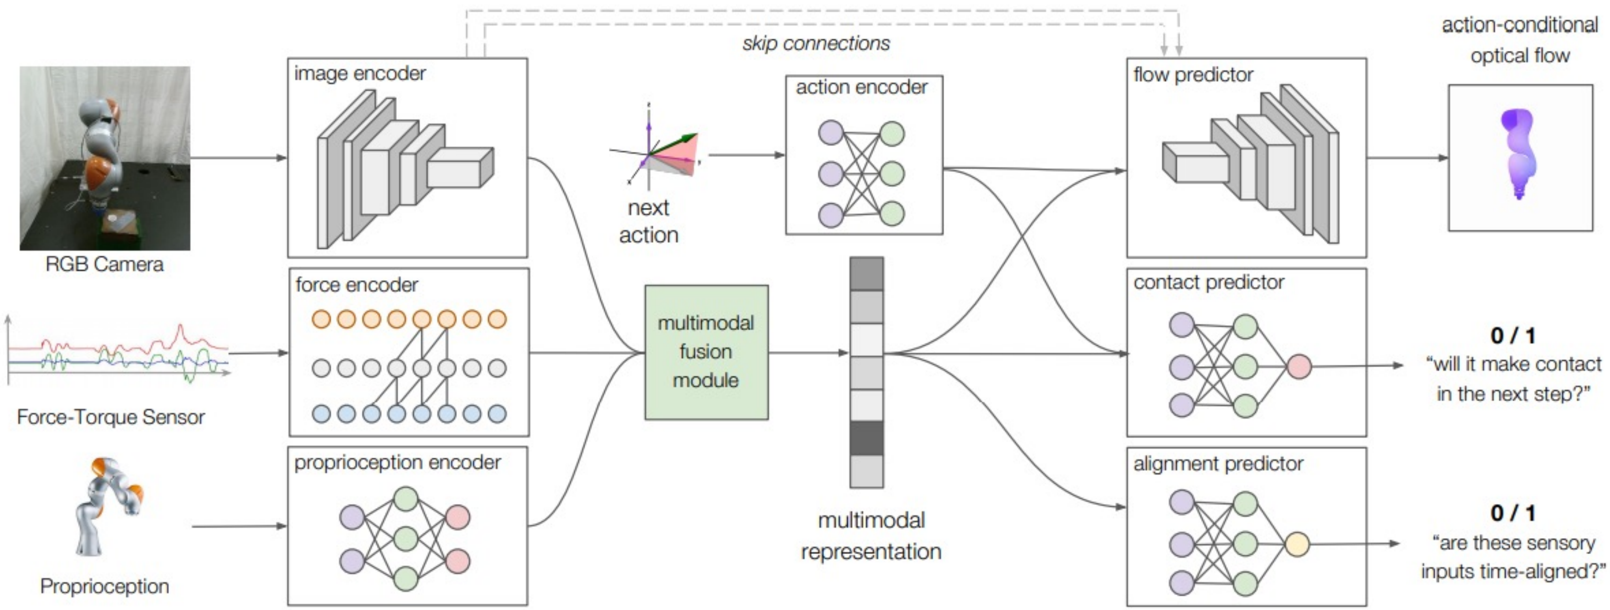

Fig. 2: Neural network architecture for multimodal representation learning with self-supervision. The network takes data from three different sensors as input: RGB images, F/T readings over a 32ms window, and end-effector position and velocity. It encodes and fuses this data into a multimodal representation based on which controllers for contact-rich manipulation can be learned. This representation learning network is trained end-to-end through self-supervision.

Domain randomization (best student paper)

Closing the Sim-to-Real Loop: Adapting Simulation Randomization with Real World Experience

Yevgen Chebotar[1,2] Ankur Handa[1] Viktor Makoviychuk[1]
Miles Macklin[1,3] Jan Issac[1] Nathan Ratliff[1] Dieter Fox[1,4]

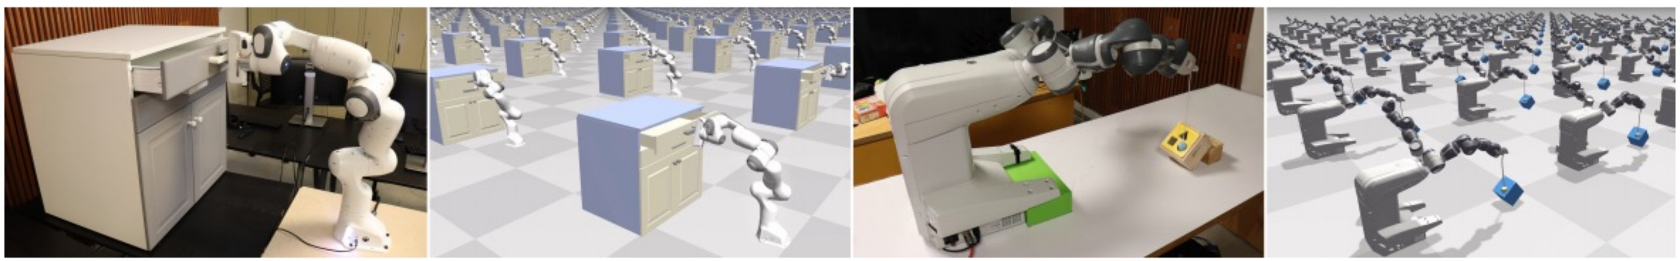

Fig. 1. Policies for opening a cabinet drawer and swing-peg-in-hole tasks trained by alternatively performing reinforcement learning with multiple agents in simulation and updating simulation parameter distribution using a few real world policy executions.

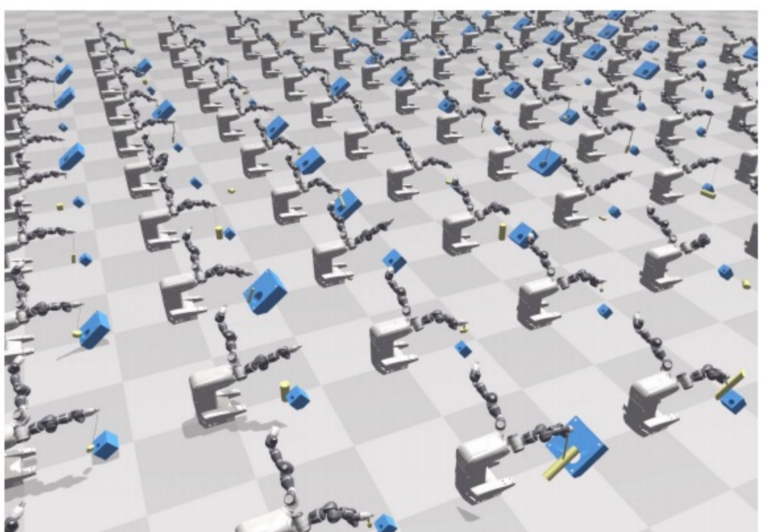

Conference on Robot Learning (CoRL) - 2019 Edition

The Conference on Robot Learning (CoRL) is a new annual international conference focusing on the intersection of robotics and machine learning. The first meeting (CoRL 2017) and the second meeting (CoRL 2018) were held in Mountain View, California on November 13 - 15, 2017 and in Zurich, Switzerland on October 29 - 31, 2018, respectively. They brought together about 350 of the best researchers working on robotics and machine learning.

CoRL 2019 will be held on October 30th- November 1st, 2019, in Osaka, Japan.

Perception and Manipulation

[1B] Perception and Manipulation (09h45 - 10h30)

Oral presentation (10 min presentation + 4 min QA)

Chair: Eiji Uchibe (Advanced Telecommunications Research Institute International)

1B-01

Towards Learning to Detect and Predict Contact Events on Vision-based Tactile Sensors

Yazhan Zhang (HKUST)*
Weihao Yuan (HKUST)
Zicheng Kan (HKUST)
Michael Yu Wang (HKUST)

1B-02

Multi-Frame GAN: Image Enhancement for Stereo Visual Odometry in Low Light

Nan Yang (Technical University of Munich)*
Eunah Jung (TUM)
Daniel Cremers (TU Munich)

1B-03

Learning to Manipulate Objects Collections Using Grounded State Representations

Matthew Wilson (University of Utah)*
Tucker Hermans (University of Utah)

Stereo VO

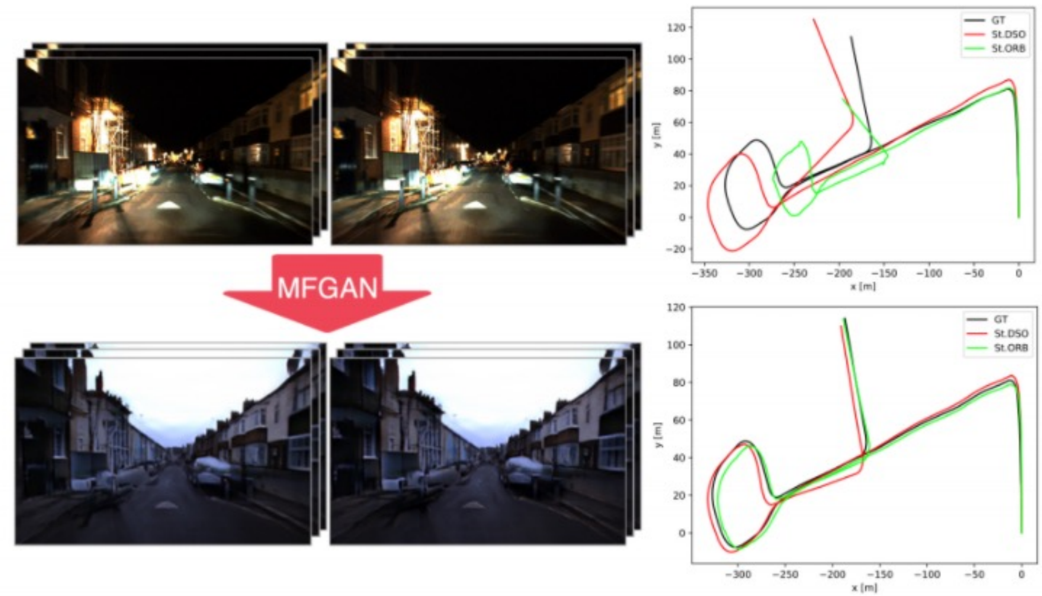

Figure 1: We propose Multi-Frame GAN (MFGAN) for stereo VO in challenging low light environment. The MFGAN takes two consecutive stereo image pairs and outputs the enhanced stereo images while preserving temporal and stereo consistency. On the right side, the estimated trajectories by the state-of-the-art stereo feature-based VO method Stereo ORB-SLAM and the state-of-the-art direct VO method Stereo DSO are presented. Due to the low image gradient, dynamic lighting and halo, Stereo DSO and Stereo ORB-SLAM cannot achieve good tracking accuracy in the night scene. With the translated images from MFGAN, the performance of both methods is notably improved.

- MF = Multi-Frame
- GAN = Generative Adversarial Networks
- VO = Visual Odometry

Manipulating Objects

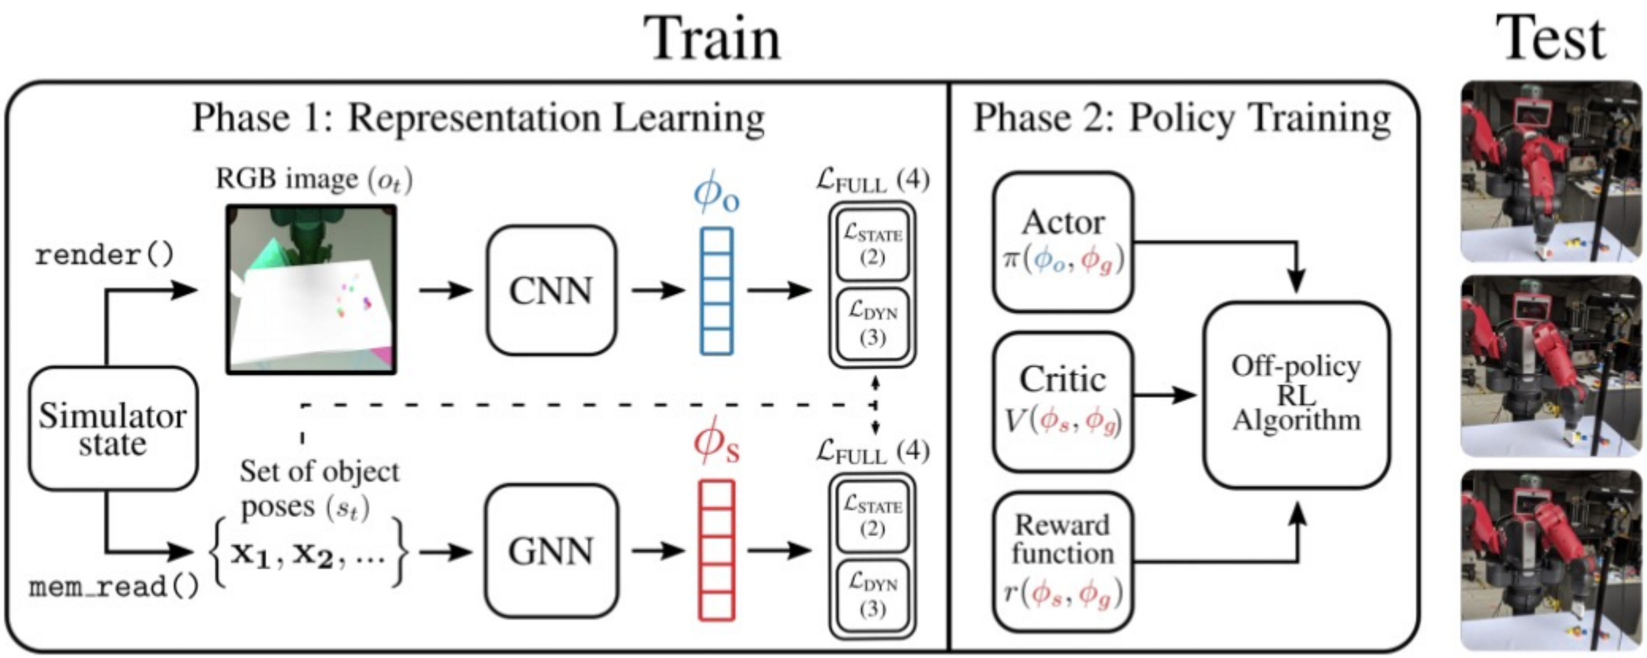

Figure 1: Cartoon diagram of our approach. We first independently train two encoder networks, one convolutional neural network (CNN) and one graph neural network (GNN) using a multi-object state and dynamics loss function. Then, during our RL phase, we embed the observation: $o \xrightarrow{\text{CNN}} \phi_o$, state: $s \xrightarrow{\text{GNN}} \phi_s$, and goal: $g \xrightarrow{\text{GNN}} \phi_g$, and we use the embeddings in an asymmetric actor critic framework [10] to train a multi-object policy π.

Planning and Control

[1F] Planning and Control (14h00 - 15h00)

Oral presentation (10 min presentation + 4 min QA)

Chair: Kostas Bekris (Rutgers University)

1F-01

Connectivity Guaranteed Multi-robot Navigation via Deep Reinforcement Learning

Juntong Lin (Sun Yat-sen University)
Xuyun Yang (Sun Yat-sen University)
Peiwei zheng (Sun Yat-sen University)
HUI CHENG (Sun Yat-Sen University)*

1F-02

Dynamics Learning with Cascaded Variational Inference for Multi-Step Manipulation

Kuan Fang (Stanford Univeristy)*
Yuke Zhu (Stanford University)
Animesh Garg (Stanford, Nvidia)
Silvio Savarese (Stanford University)
Li Fei-Fei (Stanford University & Google)

1F-03

An Online Learning Procedure for Feedback Linearization Control without Torque Measurements

Marco Capotondi (Private)
Giulio Turrisi (Sapienza, University of Rome)
Claudio Roberto Gaz (Sapienza Università di Roma)
Valerio Modugno (Sapienza, university of Rome)*
Giuseppe Oriolo (La Sapienza)
Alessandro De Luca (Sapienza University of Rome)

1F-04

Learning from My Partner's Actions: Roles in Decentralized Robot Teams

Dylan P Losey (Stanford University)*
Mengxi Li (Stanford University)
Jeannette Bohg (Stanford)
Dorsa Sadigh (Stanford)

Dynamics Learning

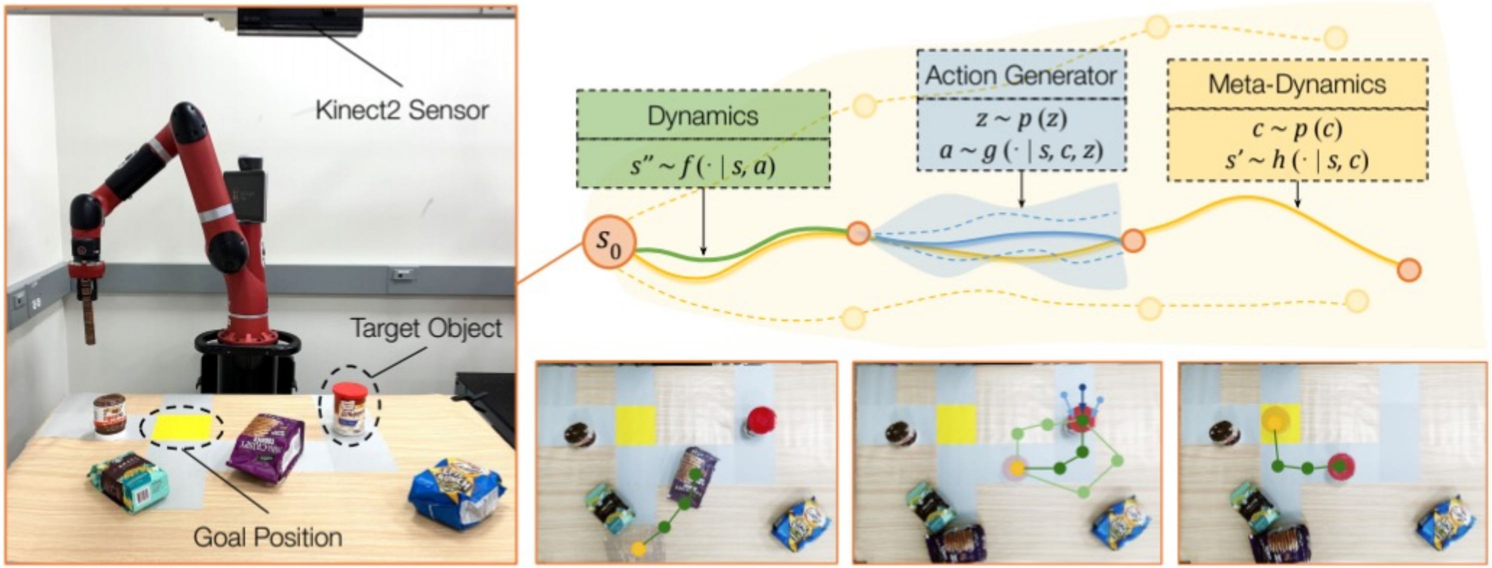

Figure 1: Hierarchical planning in latent spaces for multi-step manipulation tasks. The manipulation tasks shown in the figure requires the robot to move the target object to a goal position through specified regions (marked by grey tiles). In presence of an obstacle, the planner needs to move the obstacles aside and then move the target. We propose to use three tightly coupled modules: dynamics model, meta-dynamics model and action generator (see details in Sec. 3) to hierarchically generate plans for the task goal. Planning in learned latent spaces, our method first predicts subgoals (yellow) and then generates plausible actions (blue). The optimal plan is chosen by predicting resultant state trajectories (green) of the sampled actions. The selected plan is in darker colors.

Reinforcement Learning

[2B] Reinforcement Learning 1 (09h45 - 10h30)

Oral presentation (10 min presentation + 4 min QA)

Chair: Chelsea Finn (Stanford University)

2B-01

Worst Cases Policy Gradients

Charlie Tang (Apple Inc.)*
Jian Zhang (Apple Inc.)
Russ Salakhutdinov (University of Toronto)

2B-02

Bayesian Optimization Meets Riemannian Manifolds in Robot Learning

Noémie Jaquier (Idiap Research Institute)*
Leonel Rozo (Bosch Center for Artificial Intelligence)
Sylvain Calinon (Idiap Research Institute)
Mathias Buerger (BCAI)

2B-03

Graph Policy Gradients for Large Scale Robot Control

Arbaaz Khan (University of Pennsylvania)*
Ekaterina Tolstaya (University of Pennsylvania)
Alejandro Ribeiro (University of Pennsylvania)
Vijay Kumar (University of Pennsylvania)

Worst case RL

Figure 7: *CARLA scenarios. Left: 3D view. Right: top-down view.*

Table 3: *Collision and (success rates) for different α in CARLA scenarios.*

Unprotected Left Turn: (Town05)		
$\alpha=0.2$	$\alpha=0.5$	$\alpha=1.0$
0% (100%)	24% (76%)	42% (58%)
Merge: (Town04)		
$\alpha=0.2$	$\alpha=0.5$	$\alpha=1.0$
2% (83%)	4% (89%)	24% (76%)

Large-scale Robot Control

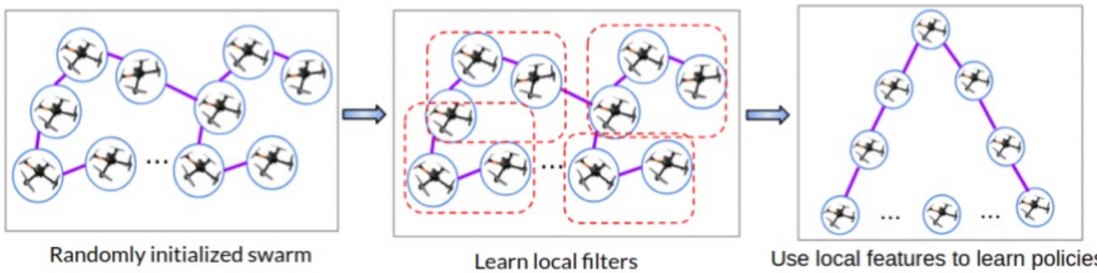

Figure 1: **Graph Policy Gradients.** Robots are randomly initialized and, based on some user set thresholds, a graph is defined. Information from K-hop neighbors is aggregated at each node by learning local filters. These local features are then used to learn policies to produce desired behavior.

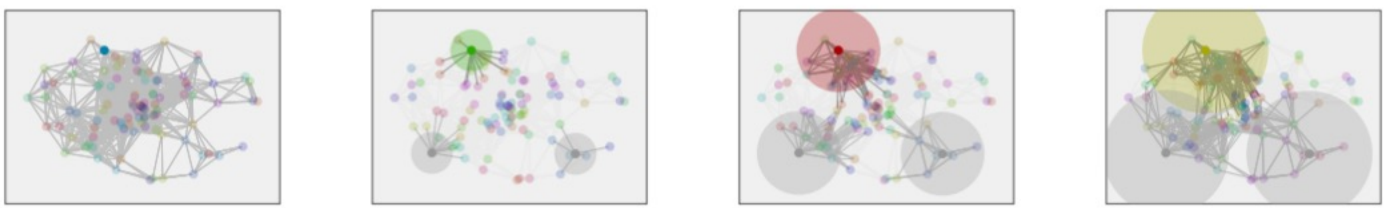

(a) $\mathbf{y}^0$ (b) $\mathbf{y}^1$ (c) $\mathbf{y}^2$ (d) $\mathbf{y}^3$

Figure 2: **Graph Convolutional Networks**. GCNs aggregate information between nodes and their neighbors. For each k-hop neighborhood (illustrated by the increasing disks), record $\mathbf{y}_{kn}$ (Eq. 3) to build $\mathbf{z}$ which exhibits a regular structure (Eq. 5). **a)** The value at each node when initialized and at the **b)** one-hop neighborhood. **c)** two-hop neighborhood. **d)** three-hop neighborhood.

Reinforcement Learning 2

[2F] Reinforcement Learning 2 (14h00 - 15h00)

Oral presentation (10 min presentation + 4 min QA)

Chair: Jens Kober (TU Delft)

2F-01
Curious iLQR: Resolving Uncertainty in Model-based RL
Sarah M.E Bechtle (Max Planck Institute for Intelligent Systems)*
Yixin Lin (Facebook AI Research)
Akshara Rai (Facebook)
Ludovic Righetti (New York University)
Franziska Meier (Facebook AI Research)

2F-02
MAT: Multi-Fingered Adaptive Tactile Grasping via Deep Reinforcement Learning
Bohan Wu (Columbia University)*
Iretiayo Akinola (Columbia University)
Jacob Varley (Google)
Peter K Allen (Columbia University)

2F-03
Adversarial Active Exploration for Inverse Dynamics Model Learning
Zhang-Wei Hong (Preferred Networks)
Tsu-Jui Fu (UC Santa Barbara)
Tzu-Yun Shann (University of British Columbia)
Yi Hsiang Chang (National Tsing Hua University)
Chun-Yi Lee (National Tsing Hua University)*

2F-04
Multi-Agent Manipulation via Locomotion using Hierarchical Sim2Real
Ofir Nachum (Google)*
Michael Ahn (Google)
Hugo Ponte (Self)
Shixiang Gu (Google Brain)
vikash kumar (Google)

Multi-agent Manipulation

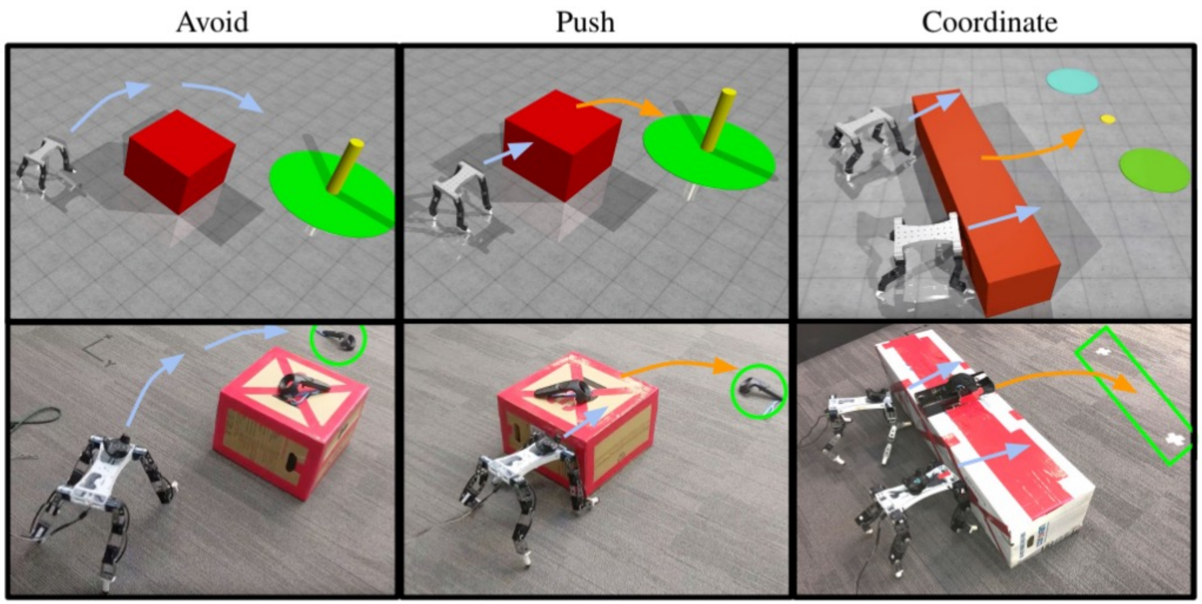

Figure 2: We consider three quadrupedal locomotion tasks of increasing complexity, utilizing the **D'Kitty** robot (see Section 4.1 for details on this robot). From left to right, we present the simulated (top row, using MuJoCo [13]) and real-world (bottom row) versions of the three tasks: *Avoid*, in which the quadruped must walk to a target location while avoiding a block object; *Push*, in which a quadruped must push a block object to a desired location; and *Coordinate*, in which two quadrupeds coordinate to push a long block to a target location and orientation. We utilize HTC Vive controllers and trackers to track the real-world position and orientation of agents, objects, and (for *Avoid* and *Push*) the desired target locations.